

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE**

In re application of: <b>Black et al.</b>	§	Group Art Unit: <b>2131</b>
	§	
Serial No.: <b>09/931,301</b>	§	Examiner: <b>Chai, Longbit</b>
	§	
Filed: <b>August 16, 2001</b>	§	
	§	
For: <b>Presentation of Correlated Events</b>	§	
<b>as Situation Classes</b>		

**Commissioner for Patents  
P.O. Box 1450  
Alexandria, VA 22313-1450**

**35525**  
PATENT TRADEMARK OFFICE  
CUSTOMER NUMBER

**APPEAL BRIEF (37 C.F.R. 41.37)**

This brief is in furtherance of the Notice of Appeal, filed in this case on November 28, 2006.

No fees are believed to be required. If, however, any fees are required, I authorize the Commissioner to charge these fees which may be required to IBM Corporation Deposit Account No. 09-0447. No extension of time is believed to be necessary. If, however, an extension of time is required, the extension is requested, and I authorize the Commissioner to charge any fees for this extension to IBM Corporation Deposit Account No. 09-0447.

### **REAL PARTY IN INTEREST**

The real party in interest in this appeal is the following party: International Business Machines Corporation of Armonk, New York.

## **RELATED APPEALS AND INTERFERENCES**

With respect to other appeals or interferences that will directly affect, or be directly affected by, or have a bearing on the Board's decision in the pending appeal, there are no such appeals or interferences.

## **STATUS OF CLAIMS**

### **A. TOTAL NUMBER OF CLAIMS IN APPLICATION**

Claims in the application are: 1-21.

### **B. STATUS OF ALL THE CLAIMS IN APPLICATION**

1. Claims canceled: NONE.
2. Claims withdrawn from consideration but not canceled: NONE.
3. Claims pending: 1-21.
4. Claims allowed: NONE.
5. Claims rejected: 1-21.
6. Claims objected to: NONE.

### **C. CLAIMS ON APPEAL**

The claims on appeal are: 1-21.

## **STATUS OF AMENDMENTS**

There are no amendments after final rejection.

## **SUMMARY OF CLAIMED SUBJECT MATTER**

### **A. CLAIMS 1, 8, and 15 - INDEPENDENT**

Independent claims 1, 8, and 15 of the present invention are directed to a method, a computer program product, and a data processing system for reporting security situations (Specification, page 10, lines 26-29), comprising the steps of logging events by storing event attributes as an event set, wherein each event set includes a source attribute, a target attribute and an event category attribute (Specification, page 12, lines 6-23; Figure 9, block 900); classifying events as groups by aggregating events with at least one attribute within the event set as an identical value (Specification, page 12, lines 6-23; block 904); calculating severity levels for the groups, wherein a severity level for a group is a function of a number of events comprising the group and values of common elements in the group (Specification, page 12, line 24 to page 13, line 7; block 908); and reporting a group from the groups to a user as a situation, if a severity level of the group exceeds a threshold value (Specification, page 13, line 7 to page 14, line 6; Figure 9, block 910).

### **B. CLAIMS 2, 9, and 16 - DEPENDENT**

Dependent claims 2, 9, and 16 of the present invention are directed to a method, a computer program product, and a data processing system wherein the severity levels are calculated based on at least one of the number of event sets within each of the groups, the source attribute of the event sets within each of the groups, the target attribute of the event sets within each of the groups, and the event category attribute of the event sets within each of the groups (Specification page 12, line 24 to page 13, line 32).

### **C. CLAIMS 7, 14, and 21 - DEPENDENT**

Dependent claims 7, 14, and 21 of the present invention are directed to a method, a computer program product, and a data processing system which further comprise aggregating a subset of the groups into a combined group (Specification page 14, lines 7-28; Figure 7, blocks 702, 704; Figure 8, blocks 802, 804, 806, 808).

## **GROUND OF REJECTION TO BE REVIEWED ON APPEAL**

### **A. GROUND OF REJECTION 1 (Claims 1-4, 8-11, and 15-18)**

Claims 1-4, 8-11, and 15-18 stand rejected under 35 U.S.C. § 102(e) as being anticipated by *Drake* et al. (U.S. Patent No. 6,347,374).

### **B. GROUND OF REJECTION 2 (Claims 5-7, 12-14, and 19-21)**

Claims 5-7, 12-14, and 19-21 stand rejected under 35 U.S.C. § 103(a) as being obvious in view of *Drake* et al. (U.S. Patent No. 6,347,374) in view of *Burrows* et al. (U.S. Patent No. 2002/0073338).

## ARGUMENT

### A. GROUND OF REJECTION 1

#### A.1. 35 U.S.C. § 102(e), Anticipation: Claims 1-4, 8-11, and 15-18

The Final Office Action rejects claims 1-4, 8-11, and 15-18 under 35 U.S.C. §102(e) as being anticipated by *Drake et al.*, Event Detection, U.S. Patent No. 6,347,374 (February 12, 2002) (hereinafter "*Drake*"). This rejection is respectfully traversed.

As to claims 1, 8, and 15, the Final Office Action states:

As per claim 1, 8 and 15, Drake teaches a method in a data processing system for reporting security situations, comprising the steps of:

logging events by storing event attributes as an event set, wherein each event set includes a source attribute, a target attribute and an event category attribute (Drake: Column 12 Line 43 – 45, Column 15 Line 60 – 67 and Column 19 Line 2 – 4 & Line 16 – 18: (a) collecting statistic data based on the following key parameters: event category, user, platform and interval (Drake: Column 15 Line 65 – 67); (b) in one embodiment, as taught by Drake, the event category can be "authentication failures" (Drake: Column 14 Line 18 – 21), which also appears on the instant application (SPEC: Page 12 Line 22 – 23); where the source attribute is considered as a specific user entity and the target attribute is considered as a specific platform – the target / platform attribute, in this case, is merely interpreted as an entity that detects / recovers the fault in an event detection system (Drake: Column 19 Line 2 – 4 & Line 16 – 18));

classifying events as groups by aggregating events with at least one attribute within the event set as an identical value (Drake: Column 7 Line 49 – 52, Column 12 Line 2 – 4, Column 12 Line 38 – 41 and Column 16 Line 1 – 8, Column 11 Line 38 – 50, Column 14 Line 18 – 21, Column 15 Line 60 – 63: (a) Drake teaches aggregating the correlated raw events into event groups with at least one attribute within the event set as an identical value such as same user ID, or same group type as "authentication failure" to generate an alert of severity situations, (b) the event detection system features by collecting statistic data based on category, user and platform in order to analyze statistical data and detect intrusion events (Drake: Column 15 Line 60 – 63), (c) dynamically increasing (i.e. aggregating) the number and type of events added to the database during a suspected intrusion attempts (Drake: Column 7 Line 49 – 52)).

calculating severity levels for the groups, wherein a severity level for a group is a function of a number of events comprising the group and values of common elements in the group (Drake: Column 12 Line 29 – 30, Column 16 Line 15 – 18, Column 11 Line 38 – 50, Column 16 Line 15 – 18 and Column 14 Line 18 – 21: the "authentication failure" is qualified to meet the severity level as an event caused by the failures of a user login – i.e. a group of events are formed and



established as "authentication failures", for example, as set forth in the discussion above and the severity level for this group is indeed a function of a number of events comprising the group and values of common elements such as common user entity (source) and platform entity (target) and the mathematical function of a number of events are compared against predetermined thresholds in order to assign a severity level to that event group).

reporting a group from the groups to a user as a situation, if a severity level of the group exceeds a threshold value (Drake: Column 11 Line 38 – 50 and Column 14 Line 18 – 21: the "authentication failure" is qualified to meet the severity level as an event caused by the failures of a user login when the aggregating events exceed the predetermined number (i.e., threshold = 3) as taught by Drake).

Final Office Action dated August 28, 2006, pages 5-6.

In addition, the Examiner further states:

1. Applicant's arguments with respect to the subject matter of the instant claims have been fully considered but are not persuasive.
2. As per claim 1, 8 and 15, Applicant first asserts: "Drake's storing events in a database table is not the same as grouping event attributes into an event set (Remarks: Page 9, 1<sup>st</sup> Para)". Examiner respectfully disagrees with the following rationale:
  - Drake teaches (a) the event detection system features by collecting statistic data based on category, user and platform in order to analyze statistical data and detect intrusion events (Drake: Column 15 Line 60 – 63); (b) dynamically increasing (i.e. aggregating) the number and type of events added to the database during a suspected intrusion attempts (Drake: Column 7 Line 49 – 52). Therefore, Examiner notes Drake does teach: "grouping event attributes into an event set" and as such Applicant's arguments are respectfully traversed.
3. Applicant first asserts: Drake does not teach: "an event set includes a source attribute, a target attribute and an event category attribute (Remarks: Page 9, 2<sup>nd</sup> Para). Examiner respectfully disagrees with the following rationale:
  - Drake teaches (a) collecting statistic data based on the following key parameters: event category, user, platform and interval (Drake: Column 15 Line 65 – 67); (b) in one embodiment, as taught by Drake, the event category can be "authentication failures" (Drake: Column 14 Line 18 – 21 ), which also appears on the instant application (SPEC: Page 12 Line 22 – 23); where the source attribute is considered as a specific user entity and the target attribute is considered as a specific platform – the target / platform attribute, in this case, is merely interpreted as an entity that detects / recovers the fault in an event detection system (Drake: Column 19 Line 2 – 4 & Line 16 – 18). Therefore, Examiner notes Drake does teach: "an event set includes a source attribute, a target attribute and an event category attribute"
  - Applicant's arguments fail to comply with 37 CFR 1.111(b) because they amount to a general allegation that the claims define a patentable invention

without specifically pointing out how the language of the claims is patentably distinguished over the prior art and as such Applicant's arguments are respectfully traversed.

4. Applicant asserts: Drake does not teach: "calculating severity levels for the groups, wherein a severity level for a group is a function of a number of events comprising the group and values of common elements in the group (Remarks: Page 9, 3<sup>rd</sup> Para). Examiner respectfully disagrees with the following rationale:

- Drake teaches calculating severity levels for the groups, wherein a severity level for a group is a function of a number of events comprising the group and values of common elements in the group (Drake: Column 12 Line 29 – 30, Column 16 Line 15 – 18, Column 11 Line 38 – 50, Column 16 Line 15 – 18 and Column 14 Line 18 – 21 : with respect to "authentication failure"). On this regard, Applicant further argues that Drake merely discloses the use of assigning a security level to a record and is not to a group of events as recited in the claim (Remarks: Page 10, 2<sup>nd</sup> Para). Examiner respectfully disagrees because a group of events are formed and established as "authentication failures", for example, as set forth in the discussion above and the severity level for this group is indeed a function of a number of events comprising the group and values of common elements such as common user entity (source) and platform entity (target) and the mathematical function of a number of events are compared against predetermined thresholds in order to assign a severity level to that event group (Drake: Column 12 Line 29 – 30, Column 16 Line 15 – 18, Column 11 Line 38 – 50, Column 16 Line 15 – 18 and Column 14 Line 18 – 21) and as such Applicant's arguments are respectfully traversed.

Final Office Action dated August 28, 2006, pages 2-4.

Independent claim 1, which is representative of independent claims 8 and 15 with regard to similarly recited subject matter, reads as follows:

1. A method in a data processing system for reporting security situations, comprising the steps of:

- logging events by storing event attributes as an event set, wherein each event set includes a source attribute, a target attribute and an event category attribute;

- classifying events as groups by aggregating events with at least one attribute within the event set as an identical value;

- calculating severity levels for the groups, wherein a severity level for a group is a function of a number of events comprising the group and values of common elements in the group; and

- reporting a group from the groups to a user as a situation, if a severity level of the group exceeds a threshold value.

A prior art reference anticipates the claimed invention under 35 U.S.C. §102 only if every element of a claimed invention is identically shown in that single reference, arranged as they are

in the claims. *In re Bond*, 910 F.2d 831, 832, 15 U.S.P.Q.2d 1566, 1567 (Fed. Cir. 1990). All limitations of the claimed invention must be considered when determining patentability. *In re Lowry*, 32 F.3d 1579, 1582, 32 U.S.P.Q.2d 1031, 1034 (Fed. Cir. 1994). Anticipation focuses on whether a claim reads on the product or process a prior art reference discloses, not on what the reference broadly teaches. *Kalman v. Kimberly-Clark Corp.*, 713 F.2d 760, 218 U.S.P.Q. 781 (Fed. Cir. 1983). In this case each and every feature of the presently claimed invention is not identically shown in the cited reference, arranged as they are in the claims.

Appellants first address the rejection of claim 1, which is representative of independent claims 8 and 15 with regard to similarly recited subject matter. Contrary to the Examiner's assertions, *Drake* does not anticipate claim 1 because *Drake* does not teach all of the features of claim 1. In particular, the Examiner states that the feature, logging events by storing event attributes as an event set, wherein each event set includes a source attribute, a target attribute and an event category attribute, is taught by the following passages of *Drake*:

The event detection system uses database table relationships to store events and to map event instance data to meta-data.

*Drake*, column 12, lines 43-45.

Statistical processors are used to process the following Event detection system features: collect statistical data by category, user, and platform; and analyze statistical data and detect events based on statistical profiles.

Implementation of statistical processing in the event detection system 10 requires design and implementation of two components. The first is a design of a database schema to permit the storage of statistical data by the following key parameters: event category; user; platform; and interval.

*Drake*, column 15, lines 60-67.

Alternatively, event detection system processes may be launched as a part of startup processes for the platform on which the processes are to be run.

*Drake*, column 19, lines 2-4.

Usually, only the process that detects the fault is sufficiently closely coupled to the fault's particular fault condition to recover.

*Drake*, column 19, lines 16-18.

*Drake* does not anticipate claim 1 because *Drake* does not teach the above recited feature of claim 1. Specifically, *Drake* does not teach storing event attributes as an event set. “The identical invention must be shown in as complete detail as is contained in the ... claim.” *Richardson v. Suzuki Motor Co.*, 868 F.2d 1226, 1236, 9 USPQ2d 1913, 1920 (Fed. Cir. 1989). *Drake* specifically states that a database table is used to store events. Contrary to the Examiner’s assumptions, storing events in a database table is not the same as grouping event attributes into an event set.

Furthermore, *arguendo*, even if *Drake* teaches an event set, *Drake* does not teach an event set that includes 1) a source attribute, 2) a target attribute, and 3) an event category attribute. *Drake* collects data by 1) category, 2) user, and 3) platform. (*Drake*, column 15, lines 60-67). *Drake*’s user attribute identifies the user performing illegal actions. (*Drake*, column 2, lines 35-40). Assuming that this attribute teaches the source attribute as recited in claim 1, *Drake* does not teach an event set that includes the target attribute. *Drake* teaches a platform attribute. *Drake* does not specifically define the term platform. The meaning of a platform to a person having ordinary skill in the art describes some sort of framework, either in hardware or software, which allows software to run. (*Wikipedia*). An example of a platform would be Windows NT. (*Drake*, column 8, lines 33-34). The target attribute as recited in claim 1 identifies the destination of the attack, i.e. the target computer(s). (*Specification*, page 11, lines 5-11). *Drake* merely comprises processes for fault management (*Drake*, column 19, lines 2-18), and does not mention anything about having the platform as a ‘target’ (e.g., intended destination) of any event (e.g., security breach). *Drake*’s platform attribute is clearly not the same as the target attribute of the presently claimed invention. Because *Drake* does not teach an event set that includes the target attribute, *Drake* does not teach or suggest the above recited feature of claim 1.

Furthermore, *Drake* does not teach the feature of calculating severity levels for the groups, wherein a severity level for a group is a function of a number of events comprising the group and values of common elements in the group as recited in claim 1. The Examiner asserts otherwise, citing to the following passages of *Drake*:

Information such as the event detection system event number and severity level are derived by this method. At any stage of event processing, meta-data may be derived.

*Drake*, column 12, lines 29-30.

The detector 32 counts events as they arrive to see if one of the thresholds is exceeded. If a threshold is exceeded, a high severity event is generated and passed to the inserter 22 for storage in a Virtual Record in the database 12.

*Drake*, column 16, lines 15-18.

In the present embodiment, there are six, standard, defined severity levels, one of which is assigned to each Virtual Record.

<u>Level</u>	<u>Meaning</u>
0	Irrelevant or undefined
1	Potentially significant event
2	Interesting event
3	Significant event
4	Warning
5	Alert

*Drake*, column 11, lines 38-50.

For example, consider a set of rules that generates an alert on three failed logins. The rules for this alert are "three failed logins, by a user, at a platform, without an intervening successful login or system restart".

*Drake*, column 14, lines 18-21.

The first passage above discloses that a rules-based processing method applied to an event record when the record is inserted into the database is used to derive an event detection system event number and severity level. The second passage discloses counting the event to see if one of the thresholds is exceeded. The third passage discloses the various severity levels, such as irrelevant, potentially significant, interesting, significant, warning, and alert, and that each record is assigned one of the severity levels. And the fourth passage discloses a rules-based alert which generates an alert based on three failed logins by a user.

However, the passages above do not teach calculating severity levels for the groups, wherein a severity level for a group is a function of a number of events comprising the group and values of common elements in the group. The passages merely disclose the use of assigning a

severity level to a record, and that the rules-based alert may be used to generate an alert upon the failure of a user's 3<sup>rd</sup> login attempt. There is no discussion in *Drake* of calculating a severity level for a group of events as recited in the claimed invention. The Examiner alleges that "authentication failure" is qualified to meet the severity level as an event by the failures of a user login. However, determining whether an alert should be generated based on multiple unsuccessful logins is not the same as calculating a severity level for a group of events. Rather, as shown above in column 11, lines 38-50 and column 12, lines 29-30, *Drake* does not teach assigning severity levels to groups of events, but rather *Drake* explicitly teaches that one severity level is assigned to each Virtual Record. As disclosed in column 6, lines 6-8, *Drake* teaches that a Virtual Record is a "standardized flat representation of an event in a normalized format". Thus, even though *Drake* derives security levels, these levels are derived for each Virtual Record, which represent a single event, rather than a group of events as recited in claim 1. In addition, *Drake* does not mention that there is a severity level calculated for the group itself. Instead, *Drake* merely discloses that an alert is generated if a specific number of the same event occurs (e.g., three failed login attempts by a particular user).

Furthermore, the passages above also do not teach that the severity level for a group is a function of a number of events comprising the group and values of common elements in the group. Thus, the common elements in the group have values which are used to calculate the severity level of the group. *Drake* does not mention a severity level calculated for the group itself and that the severity level of the calculated group is based on common elements in the group. *Drake* merely discloses that an alert is generated based on the occurrence of a specific number of events (e.g., three failed login attempts). Thus, while *Drake* may derive and assign severity levels to individual records in the database tables, *Drake* does not teach anything about calculating a severity level for a group of events, nor does *Drake* teach that the calculated group severity level is based on a number of events comprising the group and values of common elements in the group.

Because *Drake* does not disclose all the features as recited in claim 1, *Drake* does not anticipate claim 1. For the same reasoning, *Drake* does not anticipate independent claims 8 and 15. At least by virtue of their dependency on claims 1, 8, and 15, respectively, *Drake* also does not teach the features in dependent claims 2-4, 9-11, and 16-18.

Furthermore, claims 2-4, 9-11, and 16-18 recite additional subject matter not taught by *Drake*. For instance, claims 2, 9, and 16 recite that severity levels are calculated based on at least one of the number of event sets within each of the groups, the source attribute of the event sets within each of the groups, the target attribute of the event sets within each of the groups, and the event category attribute of the event sets within each of the groups. As discussed in the response to the rejection of claims 1, 8, and 15 above, the feature of calculating severity levels for an event group in these claims are neither taught nor suggested by *Drake*. Furthermore, an event group containing a target attribute is neither taught nor suggested by *Drake*. Consequently, it is respectfully urged that the rejection of claims 1-4, 8-11, and 15-18 under 35 U.S.C. §102(e) has been overcome.

Accordingly, Appellants respectfully request withdrawal of the rejection of claims 1-4, 8-11, and 15-18 under 35 U.S.C. §102(e).

## **B. GROUND OF REJECTION 2**

### **B.1. 35 U.S.C. § 103, Obviousness, Claims 5-7, 12-14, and 19-21**

The Final Office Action rejects claims 5-7, 12-14, and 19-21 under 35 U.S.C. §103 as being obvious in view of *Drake* in view of *Burrows* et al., Method and System for Limiting the Impact of Undesirable Behavior of Computers on a Shared Data Network, U.S. Patent Publication No. 2002/0073338 (June 13, 2002) (hereinafter “*Burrows*”). This rejection is respectfully traversed.

As to claims 5-7, 12-14, and 19-21, the Examiner states:

As per claim 7, 14 and 21, Drake does not disclose expressly aggregating a subset of the groups into a combined group.

Burrows teaches aggregating a subset of the groups into a combined group (Burrows: Para [0050] and Para [0046] Line 10 – 11: similar to the Figure 8 / Element 804/806/802 of the instant application, the single source computer that causes broadcast storms to any of the unspecified destination computers, as taught by Burrows, does indeed generate a combined group of events. Likewise, it applies to the similar situation of denial-of-service attacks).

It would have been obvious to a person of ordinary skill in the art at the time the invention was made to combine the teaching of Burrows within the system of Drake because (a) Drake teaches improving network security by providing an effective event detecting systems (Drake, see example, Column 2 Line 4 – 8 and Column 3 Line 34 – 35) and (b) Burrows teaches managing and tracking computer security incidents that may occur in a network computer system by effectively detecting any types of behavior and undesirable patterns of packet traffic (Burrows: Para [0019]).

As per claim 5, 12 and 19, Drake does not disclose expressly the target attribute represents one of a computer and a collection of computers.

Burrows teaches the target attribute represents one of a computer and a collection of computers (Burrows: Para [0050] and Para [0046] Line 10 – 11: the target attribute could be the single server computer that causes denial-of-service or a collection of computers such as broadcast storms).

Same rationale of combination applies herein as above in rejecting the claim 7.

As per claim 6, 13 and 20, Drake does not disclose expressly the source attribute represents one of a computer and a collection of computers.

Burrows teaches the source attribute represents one of a computer and a collection of computers (Burrows: Para [0050] and Para [0046] Line 10 – 11: the source attribute could be the single source computer that causes broadcast storms to any of the unspecified destination computers or as a collection of computers such as denial-of-service attacks).



Same rationale of combination applies herein as above in rejecting the claim 7.

Final Office Action dated August 28, 2006, pages 8-9.

Regarding claims 5-7, 12-14, and 19-21, the Examiner has failed to state a *prima facie* obviousness rejection because the proposed combination does not teach or suggest all of the features of claims 5-7, 12-14, and 19-21. A *prima facie* case of obviousness is established when the teachings of the prior art itself suggest the claimed subject matter to a person of ordinary skill in the art. *In re Bell*, 991 F.2d 781, 783, 26 U.S.P.Q.2d 1529, 1531 (Fed. Cir. 1993). All limitations of the claimed invention must be considered when determining patentability. *In re Lowry*, 32 F.3d 1579, 1582, 32 U.S.P.Q.2d 1031, 1034 (Fed. Cir. 1994). If an independent claim is nonobvious under 35 U.S.C. §103, then any claim depending therefrom is nonobvious. *In re Fine*, 837 F.2d 1071, 5 USPQ2d 1596 (Fed. Cir. 1988). In the case at hand, not all of the features of the claimed invention have been properly considered and the teachings of the references themselves do not teach or suggest the claimed subject matter to a person of ordinary skill in the art.

Addressing the rejection of claims 5-7, 12-14, and 19-21, the Examiner has failed to state a *prima facie* obviousness rejection because the combination of *Drake* and *Burrows* fails to teach or suggest all features of claim 1, 8, and 15 from which claims 5-7, 12-14, and 19-21 depend, respectively. As discussed above, *Drake* does not teach the claimed feature of logging events by storing event attributes as an event set, wherein each event set includes a source attribute, a target attribute and an event category attribute. Because *Drake* specifically teaches collecting data by 1) category, 2) user, and 3) platform, *Drake* also does not suggest the features of claim 1 wherein the event set includes a target attribute. Furthermore, *Burrows* does not teach or suggest all of the features of claim 1. *Burrows* is directed to detecting undesirable behavior in a network system. Specifically, *Burrows* envisions using a packet traffic monitor to determine the existence and source of any pattern of undesirable behavior. (*Burrows*, paragraph 40.) However, *Burrows* does not teach or suggest logging events by storing event attributes as an event set, wherein each event set includes a source attribute, a target attribute and an event category attribute.

Because neither *Drake* nor *Burrows* teaches or suggests all of the features of claim 1, 8, and 15 and because claims 5-7, 12-14, and 19-21 depend from claims 1, 8, and 15, respectively, the proposed combination of *Drake* and *Burrows* when considered as a whole does not teach or

suggest all of the features of claims 5-7, 12-14, and 19-21. Accordingly, the Examiner has failed to state a *prima facie* obviousness rejection of claims 5-7, 12-14, and 19-21.

Furthermore, the proposed combination of *Drake* and *Burrows*, when considered as a whole, also does not teach or suggest all of the additional features of claims 5-7, 12-14, and 19-21. For instance, claims 7, 14, and 21 recite “aggregating a subset of the event groups into a combined group.” The Examiner points to the following passages in *Burrows* as teaching this feature:

In one embodiment, the packet traffic monitor observes the network and thereby detects and localizes all broadcast packets traffic. Observing more than a predetermined number of broadcast packets within a predetermined time period implies that a broadcast storm is underway. It is likely that the packet is correctly addressed, and that knowing the source MAC address and the network topology will point to a particular port of a forwarding device, e.g., switch port, to be disabled. In another embodiment, the per-port broadcast ingress packet counters can be used to trace broadcast packets to their source. This approach is used if the packet traffic monitor fails at determining the source, possibly because of incorrectly formatted packets or because the misbehaving host has not been seen on the network before (unknown MAC address). This detection approach is less timely than the prior approach since the process of retrieving these counters from the switch is extensive and it cannot be executed often.

*Burrows*, para [0050].

For example, the monitor can detect too many packets destined to an overloaded server, too many probe packets directed to a firewall or too many ARP request packets.

*Burrows*, para [0046], lines 10-11.

The cited passages of *Burrows* do not mention aggregating a subset of an event group into a combined event group, as recited in claims 7, 14, and 21. The first passage discusses determining the source of a broadcast storm. A broadcast storm occurs when a host emits a continuous stream of broadcast packets. The second passage merely states that the monitor can detect too many packets destined to an overloaded server. Neither passage teaches or suggests aggregating a subset of an event group into a combined event group, as recited in claims 7, 14, and 21.

Even if the missing elements of the rejected claims existed in the prior art, for the rejected claims to be obvious there must be some motivation or incentive from the prior art to modify or

combine the reference teachings to achieve the present invention. The Examiner does not provide any motivation from either reference that making all the necessary modifications to the reference teachings to achieve the present invention would be desirable. If the Examiner cannot make such a showing, then the Examiner has simply relied on hindsight with the benefit of Appellants' disclosure to develop an incentive for the changes, which in fact, would not be obvious to one of ordinary skill in the art at the time the invention was made. Therefore, the rejection of claims 5-7, 12-14, and 19-21 under 35 U.S.C. §103(a) has been overcome.

Accordingly, Appellants respectfully request the withdrawal of rejection of claims 5-7, 12-14, and 19-21 under 35 U.S.C. §103(a).

/Cathrine K. Kinslow/  
Cathrine K. Kinslow  
Reg. No. 51,886  
**YEE & ASSOCIATES, P.C.**  
PO Box 802333  
Dallas, TX 75380  
(972) 385-8777

## **CLAIMS APPENDIX**

The text of the claims involved in the appeal are:

1. A method in a data processing system for reporting security situations, comprising the steps of:

logging events by storing event attributes as an event set, wherein each event set includes a source attribute, a target attribute and an event category attribute;

classifying events as groups by aggregating events with at least one attribute within the event set as an identical value;

calculating severity levels for the groups, wherein a severity level for a group is a function of a number of events comprising the group and values of common elements in the group; and

reporting a group from the groups to a user as a situation, if a severity level of the group exceeds a threshold value.

2. The method of claim 1, wherein the severity levels are calculated based on at least one of the number of event sets within each of the groups, the source attribute of the event sets within each of the groups, the target attribute of the event sets within each of the groups, and the event category attribute of the event sets within each of the groups.

3. The method of claim 1, wherein the events include at least one of a web server event, an electronic mail event, a Trojan horse, denial of service, a virus, a network event, an authentication failure, and an access violation.

4. The method of claim 1, further comprising:

calculating the threshold value based on at least one of the source attribute of the event sets within the group, the target attribute of the event sets within the group, the event category attribute in each event set of the group, and the number of attributes in each event set of the group that are held constant across all of the event sets in the group.

5. The method of claim 1, wherein the target attribute represents one of a computer or a collection of computers.

6. The method of claim 1, wherein the source attribute represents one of a computer or a collection of computers.

7. The method of claim 1, further comprising:

aggregating a subset of the groups into a combined group.

8. A computer program product in a computer readable medium for reporting security events, comprising instructions for:

logging events by storing event attributes as an event set, wherein each event set includes a source attribute, a target attribute and an event category attribute;

classifying events as groups by aggregating events with at least one attribute within the event set as an identical value;

calculating severity levels for the groups, wherein a severity level for a group is a function of a number of events comprising the group and values of common elements in the group; and reporting a group from the groups to a user as a situation, if a severity level of the group exceeds a threshold value.

9. The computer program product of claim 8, wherein the severity levels are calculated based on at least one of the number of event sets within each of the groups, the source attribute of the event sets within each of the groups, the target attribute of the event sets within each of the groups, and the event category attribute of the event sets within each of the groups.

10. The computer program product of claim 8, wherein the events include at least one of a web server event, an electronic mail event, a Trojan horse, denial of service, a virus, a network event, an authentication failure, and an access violation.

11. The computer program product of claim 8, comprising additional instructions for: calculating the threshold value based on at least one of the source attribute of the event sets within the group, the target attribute of the event sets within the group, the event category attribute in each event set of the group, and the number of attributes in each event set of the group that are held constant across all of the event sets in the group.

12. The computer program product of claim 8, wherein the target attribute represents one of a computer or a collection of computers.

13. The computer program product of claim 8, wherein the source attribute represents one of a computer or a collection of computers.
14. The computer program product of claim 8, comprising additional instructions for:  
aggregating a subset of the groups into a combined group.
15. A data processing system for reporting security events, comprising:  
a bus system;  
a memory;  
a processing unit, wherein the processing unit includes at least one processor; and  
a set of instructions within the memory, wherein the processing unit executes the set of instructions to perform the acts of:  
logging events by storing event attributes as an event set, wherein each event set includes a source attribute, a target attribute and an event category attribute;  
classifying events as groups by aggregating events with at least one attribute within the event set as an identical value;  
calculating severity levels for the groups, wherein a severity level for a group is a function of a number of events comprising the group and values of common elements in the group; and  
reporting a group from the groups to a user as a situation, if a severity level of the group exceeds a threshold value.

16. The data processing system of claim 15, wherein the severity levels are calculated based on at least one of the number of event sets within each of the groups, the source attribute of the event sets within each of the groups, the target attribute of the event sets within each of the groups, and the event category attribute of the event sets within each of the groups.

17. The data processing system of claim 15, wherein the events include at least one of a web server event, an electronic mail event, a Trojan horse, denial of service, a virus, a network event, an authentication failure, and an access violation.

18. The data processing system of claim 15, wherein the processing unit executes the set of instructions to perform the act of:

calculating the threshold value based on at least one of the source attribute of the event sets within the group, the target attribute of the event sets within the group, the event category attribute in each event set of the group, and the number of attributes in each event set of the group that are held constant across all of the event sets in the group.

19. The data processing system of claim 15, wherein the target attribute represents one of a computer or a collection of computers.

20. The data processing system of claim 15, wherein the source attribute represents one of a computer or a collection of computers.



21. The data processing system of claim 15, wherein the processing unit executes the set of instructions to perform the act of:

aggregating a subset of the groups into a combined group.

## **EVIDENCE APPENDIX**

There is no evidence to be presented.

## **RELATED PROCEEDINGS APPENDIX**

There are no related proceedings.